Capacitated Network Design Games

Michal Feldman[1] and Tom Ron[2]

[1] Hebrew University of Jerusalem and Harvard University
mfeldman@seas.harvard.edu *
[2] Hebrew University of Jerusalem rontom@gmail.com

Abstract. We study a *capacitated* symmetric network design game, where each of n agents wishes to construct a path from a network's source to its sink, and the cost of each edge is shared equally among its agents. The uncapacitated version of this problem has been introduced by Anshelevich *et al.* (2003) and has been extensively studied. We find that the consideration of edge capacities entails a significant effect on the quality of the obtained Nash equilibria (NE), under both the utilitarian and the egalitarian objective functions, as well as on the convergence rate to an equilibrium. The following results are established. First, we provide bounds for the price of anarchy (PoA) and the price of stability (PoS) measures with respect to the utilitarian (i.e., sum of costs) and egalitarian (i.e., maximum cost) objective functions. Our main result here is that, unlike the uncapacitated version, the network topology is a crucial factor in the quality of NE. Specifically, a network topology has a bounded PoA if and only if it is *series-parallel* (SP). Second, we show that the convergence rate of best-response dynamics (BRD) may be super linear (in the number of agents). This is in contrast to the uncapacitated version, where convergence is guaranteed within at most n iterations.

1 Introduction

The construction of large networks by strategic agents has been widely studied from a game-theoretic perspective in the last decade [3, 8, 9, 25]. For a motivating example, consider the construction and maintenance of large computer networks by independent economic agents with different, and often competing, self-interests. The game-theoretic perspective offers tools and insights that are fundamental to the understanding and analysis of these settings.

In a symmetric network design game, a network is given, where each edge is associated with some cost; and a set of n agents wish to buy some path from

* This work was partially supported by the Israel Science Foundation (grant number 1219/09), by the Leon Recanati Fund of the Jerusalem School of Business Administration, the Google Inter-university center for Electronic Markets and Auctions, and the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 274919. The authors wish to thank Eli Ben-Sasson and Irit Dinur for helpful discussion.

the network's source (s) to its sink (t). Every agent chooses an s-t path, and the cost of every edge is divided equally among the agents who use it. This is often called a *fair cost-sharing* method. The game theoretic twist is the assumption that each agent chooses its path strategically, so as to minimize its cost. It is well known that Nash equilibria of this game need not be efficient, where efficiency is usually defined with respect to either the sum of the agents' costs (referred to as the *utilitarian* or *sum-cost* objective) or to the maximum cost of any agent (referred to as the *egalitarian* or *max-cost* objective).

The efficiency loss is commonly quantified using the price of anarchy (PoA) [17, 23] and price of stability (PoS) [3] measures; the former refers to the ratio between the cost of the worst Nash equilibrium and the social optimum, whereas the latter refers to the ratio between the cost of the best Nash equilibrium and the social optimum. The network design game described above is fairly easy to analyze. The PoA is known to be tightly bounded by n with respect to the utilitarian objective function[3] [3]. It is not too difficult to see that the same bound holds with respect to the egalitarian objective. In addition, the PoA is independent of the network topology, as the worst case is obtained for two parallel links. The PoS, in contrast, is always equal to 1 (with respect to both objective functions), since in a symmetric network, the profile in which all agents share the shortest path from s to t is a Nash equilibrium. Finally, best-response dynamics (i.e., dynamics in which agents sequentially apply their best-response moves) exhibits a simple structure, where convergence to a NE is guaranteed within at most n steps.

Interestingly, as we shall soon see, a lot of the aforementioned results should be attributed to the assumption that the network edges are *uncapacitated*; i.e., it is assumed that edges may hold any number of agents. While this assumption has been employed by most of the studies on strategic network formation games, we claim that in real-life applications network links have a limit on the number of agents they can serve. To reflect this observation, we introduce *capacitated* network design games, in which every edge, in addition to its cost, is also associated with a *capacity* that specifies the number of agents it can hold. We study the quality of NE in these games (using both PoA and PoS measures) and the convergence rate of best-response dynamics. We are particularly interested in the effect of the *topology* of the underlying network on the obtained results.

In cases where edges are associated with capacities, a *feasibility* problem arises (i.e., whether there exists a solution that accommodates all the agents). However, as already hinted at by [3], if a feasible solution exists, the arguments used in the uncapacitated version can be applied to show that a pure NE exists and, moreover, every best-response dynamics converges to a pure NE. This observation motivates our study.

Our contribution. For the PoA, the lower bound of n trivially carries over to the capacitated version; thus, one cannot expect for a bound better than n. The upper bound of n, however, does not carry over. In particular, we

[3] While [3] consider an underlying directed graph, this bound carries over to the undirected case.

demonstrate that the PoA can be arbitrarily high. As it turns out, however, the network topology plays a major role in the obtained PoA. A symmetric network topology G is said to be *PoA bounded* if for every symmetric network design game that is played on G, the PoA is bounded by n, independent of the edge costs and capacities. Our main result here is a full characterization of PoA-bounded network topologies. Specifically, we show that a symmetric network topology is PoA bounded if and only if it is a *series-parallel* (SP) network; i.e., a network that is built inductively by series and parallel compositions of SP networks. This result holds with respect to both the sum-cost and max-cost objectives. Moreover, for parallel-link networks, we show that the PoA (with respect to both the sum-cost and max-cost objectives) is essentially bounded by the maximum edge capacity in the network, and this is tight.

This separation between the graph topology and the assignment of edge costs and capacities reflects a separation between the underlying infrastructure and the edge characteristics. While the infrastructure is often stable over time, the edge characteristics may be modified over short time periods. A PoA bounded topology ensures that, no matter how edge characteristics evolve, the cost of a NE will never exceed n. Such topologies should be desired by network designers, who wish to guarantee the efficiency in their network despite the fact they do not control the actions of the individual users. Notably, within the class of SP networks, the worst case is obtained already for parallel links.

In contrast to the PoA, the PoS with respect to the sum-cost objective is not affected by the network topology. In particular, we provide a lower bound of $H(n)$ (i.e., the harmonic nth number) for the PoS on parallel-link graphs, and show that for every symmetric network the PoS is upper bounded by $H(n)$.

As for the max-cost objective function, for SP graphs the upper bound of n that is established for the PoA trivially carries over to the PoS, and a matching lower bound is established. For general graphs, we establish an upper bound of $n \log n$. Closing the gap between n and $\log n$ for the PoS in general graphs remains an open problem.

Most of our results for the PoA and PoS bounds are summarized in Table 1, where they are also contrasted with the corresponding results in the uncapacitated version (specified in brackets). These results suggest that the departure from the classic assumption of uncapacitated edges brings in significant differences in the quality of equilibria.

Additionally, we study the convergence rate of best-response dynamics (BRD) to a NE. Here too, the consideration of capacities introduces additional complexity that reveals itself through a slower conversion rate. While BRD in the uncapacitated version is guaranteed to converge within at most n iterations, we establish a lower bound of $\Omega(n^{3/2})$ for convergence in capacitated games. Moreover, this lower bound is obtained already in the simplest graphs; i.e., graphs that are composed of parallel links.

Finally, we note that while the feasibility problem in capacitated games is equivalent to a maximum flow computation, and thus can be solved in polynomial

		Parallel links	SP	General
sum-cost (sc)	PoA	n (n)	n (n)	unbounded (n)
	PoS	$\log n(1)$	$\log n(1)$	$\log n(1)$
max-cost (mc)	PoA	n (n)	n (n)	unbounded (n)
	PoS	n (1)	n (1)	$n \log n(1)$

Table 1. Summary of our results. The values in brackets correspond to the bounds for uncapacitated games. All the results, except for the PoS w.r.t max-cost for general networks are tight.

time, the optimization version of the problem is NP-complete (this can be easily verified through a reduction from 0-1 knapsack [14]).

Related work. Various models of network design and formation games have been extensively studied in the last decade from a game-theoretic perspective [3–5, 18, 6, 7], with a great emphasis on the PoA and PoS measures. The PoA in network design games has been also studied with respect to the *strong equilibrium* solution concept by Epstein *et al.* [8], Andelman *et al.* [2] and Albers [1].

The role that network topology plays in game-theoretic settings has been studied in various models. In the model of network routing, it has been shown by Roughgarden and Tardos [25] that the PoA is independent of the network topology. In contrast, the network topology seems to matter a lot in other settings. Some prominent examples include the following. Milchtaich [21] showed that the *Pareto efficiency* of equilibria in network routing games (with a continuum of agents) strongly depends on the network topology. In addition, topological characterizations for symmetric network games have been also provided for other equilibrium properties, including (Nash and strong) equilibrium existence (see Milchtaich [20], Epstein *et al.* [8, 9], and Holzman and Law-Yone [15, 16]), and equilibrium uniqueness (see Milchtaich [19]).

Best-response dynamics (BRD) and its convergence rate has been the subject of intensive research recently. Since every congestion game is a potential game [24, 22], BRD always converge to a pure NE. However, they may in general take exponential number of steps depending on the number of agents, as established by Fabrikant *et al.* [11]. Anshelevich *et al.* [3] established that BRD may take exponential number of steps to converge in network design games, but is polynomial for the special case of two agents. Notably, as shall be discussed in Sect. 5, the exponential convergence rate does not apply in our setting. BRD convergence has been also studied in scheduling and routing games (see Even-Dar *et al.* [10], Fotakis [13], and Feldman and Tamir [12]).

2 Model and Preliminaries

2.1 Capacitated symmetric cost sharing games

A capacitated, symmetric cost-sharing connection (CCS) game (also known as single commodity) is a tuple

$$\Delta = \langle n, G = (V, E), s, t, \{p_e\}_{e \in E}, \{c_e\}_{e \in E} \rangle,$$

where n is the number of agents and $G = (V, E)$ is an undirected graph, with $s, t \in V$ as its *source* and *sink* nodes, respectively. Every edge $e \in E$ is associated with a cost $p_e \in R^{\geq 0}$ and a capacity $c_e \in N$, where an edge capacity specifies the maximum number of agents that can use it. The set of agents $\{1, \ldots, n\}$ is also denoted by $[n]$. Every agent i wishes to construct an s-t path in G. The strategy space of an agent i, denoted Σ_i, is the set of $s - t$ paths in G, and a strategy of an agent i is denoted by $S_i \in \Sigma_i$. Since this is a symmetric game, all agents have the same strategy space. The joint action space is denoted by Σ.

We consider the *fair* cost-sharing game, where an edge's cost is shared equally by all the agents that use it in their path. Given a strategy profile $S = (S_1, \ldots, S_n)$, we denote by $x_e(S)$ the number of agents that use edge e in their path; i.e., $x_e(S) = |\{i : e \in S_i\}|$. A profile S is said to be *feasible* if for every $e \in E$, $x_e(S) \leq c_e$. The cost of agent i in a profile S is defined as

$$p_i(S) = \begin{cases} \sum_{e \in S_i} \frac{p_e}{x_e(S)} & \text{, if } S \text{ is feasible} \\ \infty & \text{, otherwise} \end{cases} \tag{1}$$

A profile S is said to be a *Nash equilibrium* if no agent can improve its cost by a unilateral deviation; i.e., for every $i, S'_i \in \Sigma_i, S_{-i} \in \Sigma_{-i}$, it holds that $p_i(S) \leq p_i(S'_i, S_{-i})$, where S_{-i} denotes the joint action of all agents except i.

Given a game Δ, let $\tau(\Delta)$ denote the set of all feasible profiles in Δ. A CCS game Δ is said to be feasible if it admits a feasible profile; i.e., $\tau(\Delta) \neq \emptyset$.

We consider two social cost functions. The *sum-cost* of a profile S is the total cost of the agents in S (and also equals the total cost of the purchased edges in S), and is given by

$$sc_\Delta(S) = \begin{cases} \sum_i p_i(S) & \text{, if } S \text{ is feasible} \\ \infty & \text{, otherwise} \end{cases}$$

The *max-cost* of a profile S is the maximum cost of any agent in S, and is given by

$$mc_\Delta(S) = \begin{cases} \max_{i \in [n]} p_i(S) & \text{, if } S \text{ is feasible} \\ \infty & \text{, otherwise} \end{cases}$$

We denote by $OPT_{sc}(\Delta)$ and $OPT_{mc}(\Delta)$ the optimal profiles with respect to the sum-cost and max-cost objectives, respectively. When clear in the context, we omit Δ, and also abuse notation and use $OPT_{sc}(\Delta)$ and $OPT_{mc}(\Delta)$ to denote the cost of the respective optimal solutions.

In the figures of the paper, every edge is associated with a tuple (c_e, p_e), denoting its capacity and cost, respectively.

2.2 Nash equilibrium existence

An uncapacitated fair cost sharing game is known to be a *potential game* [3]. Every potential game admits a pure NE [22]. Moreover, BRD (where agents sequentially apply their best-response moves) always converge to a pure NE. Capacitated versions are not guaranteed to admit a feasible solution; however, if a feasible solution exists, then so does a pure NE.

Observation 1. *[3] Let Δ be a CCS game s.t. $\tau(\Delta) \neq \emptyset$. Then, Δ admits a pure NE and every best response dynamics convergence to a NE.*

This proof relies on the existence of a potential function, $\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{x_e(S)}$, that emulates the cost of an agent when deviating from a feasible solution to another.

2.3 Efficiency loss

To quantify the efficiency loss due to strategic behavior, we use the PoA and PoS measures. The PoA is the ratio of the worst Nash equilibrium and the social optimum, and is given by $PoA_{sc}(\Delta) = \frac{\max_{S \in NE(\Delta)} sc_{\Delta}(S)}{OPT_{sc}(\Delta)}$ and $PoA_{mc}(\Delta) = \frac{\max_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)}$ with respect to the sum-cost and max-cost objectives, respectively, where $NE(\Delta)$ denotes the set of NE of Δ, and it is assumed that $NE(\Delta) \neq \emptyset$. Similarly, the PoS of sum-cost and max-cost are given by $PoS_{sc}(\Delta) = \frac{\min_{S \in NE(\Delta)} sc_{\Delta}(S)}{OPT_{sc}(\Delta)}$ and $PoS_{mc}(\Delta) = \frac{\min_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)}$, respectively.

2.4 Graph theoretic preliminaries

In this section we provide some preliminaries regarding network topologies. A *symmetric* network is an undirected graph G along with two distinguished nodes, a source s and a sink t. When clear in the context, we refer to G as the symmetric network. A CCS game is symmetric (also called single-commodity) if its underlying network is symmetric with source s and sink t, and nodes s and t are the respective source and sink of all the agents.A symmetric network G is *embedded* in a symmetric network G' if G' is isomorphic to G or to a network derived from G by applying the following operations any number of times in any order: (i) *Subdivision* of an edge (i.e., its replacement by a path of edges), (ii) *Addition* of a new edge joining two existing nodes, (iii) *Extension* of the source or the sink (i.e., addition of a new edge joining s or t with a new node, which becomes the new source or sink, respectively).

Next, we define the following operations on symmetric networks:

Identification: The *identification* operation is the collapse of two nodes into one. More formally, given a graph $G = (V, E)$ we define the *identification* of nodes $v_1 \in V$ and $v_2 \in V$ forming a new edge $v \in V$ as creating a new graph $G' = (V', E')$ where $V' = V \setminus \{v_1, v_2\} \cup \{v\}$ and E' includes the edges of E where the edges of v_1 and v_2 are now connected to v.

Parallel composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, with sources $s_1 \in V_1$ and $s_2 \in V_2$ and sinks $t_1 \in V_1$ and $t_2 \in V_2$, respectively, we define a new symmetric network $G = G_1||G_2$ as follows. Let $G' = (V_1 \cup V_2, E_1 \cup E_2)$ be the union of network. To generate $G = G_1||G_2$ we identify the sources s_1 and s_2, forming a new source node s, and identify the the sinks t_1 and t_2, forming a new sink t.
Series composition: Given two symmetric networks, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, with sources $s_1 \in V_1$ and $s_2 \in V_2$ and sinks $t_1 \in V_1$ and $t_2 \in V_2$, respectively, we define a new symmetric network $G = G_1 \rightarrow G_2$ as follows. Let $G' = (V_1 \cup V_2, E_1 \cup E_2)$ be the union network. To generate $G = G_1 \rightarrow G_2$ from G' we identify the vertices t_1 and s_2, forming a new vertex u. The network G has a source $s = s_1$ and a sink $t = t_2$.

A **series-parallel (SP)** network is a symmetric network that is constructed inductively from two SP networks by either a series composition or a parallel composition, where a single edge serves as the base of the induction. That is, a symmetric network consisting of a single edge is an SP network. In addition, given two SP networks, G_1 and G_2, the networks $G = G_1||G_2$ and $G = G_1 \rightarrow G_2$ are SP networks.

3 The sum-cost objective function

3.1 Price of anarchy (PoA)

Throughout this section, we write PoA to denote PoA_{sc} for simplicity. In uncapacitated cost sharing games, the PoA is n (tightly). This is, however, not the case in capacitated games, as demonstrated by the following proposition.

Proposition 1. *The price of anarchy with respect to the sum-cost function in CCS games can be arbitrarily high.*

Proof. Consider a CCS game with two agents and an underlying graph as depicted in Fig. 3.1(a), and suppose that y is arbitrarily larger than x. The optimal profile is where one agent uses the path s-a-t and the other uses the path s-b-t, resulting in a total cost of $4x$. However, there is a NE in which one agent uses the path s-a-b-t and the other uses the path s-b-a-t, resulting in a total cost of $4x + y$. Therefore, $PoA_{sc}(\Delta) = \frac{4x+y}{4x}$, which can be arbitrarily high.

Our goal is to characterize network topologies in which such a "bad" example cannot occur; i.e., topologies in which the PoA is always bounded, independent of the specific edge costs and capacities. The lower bound of n for a network with two parallel links motivates the following definition.

Definition 1. *A symmetric network $G = (V, E)$ with source s and sink t is PoA bounded for a family of symmetric CCS games $\mathcal{F}$ if for every symmetric CCS game $\Delta \in \mathcal{F}$ on the symmetric network G, it holds that $PoA(\Delta) \leq n$.*

Our main result is a full characterization of PoA bounded network topologies.

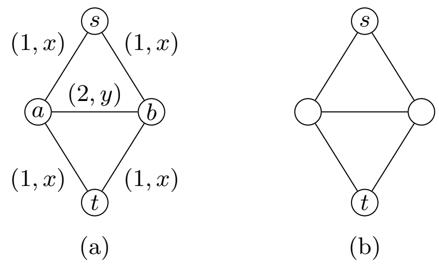

Fig. 1. (a) An example where the PoA can be arbitrarily high. (b) A *Braess* Graph.

Theorem 1. *For symmetric CCS games, a symmetric network topology G is PoA bounded w.r.t. sum-cost if and only if G is a series-parallel (SP) network.*

The proof of our characterization is composed of two parts. First, we show that for every symmetric CCS game that is played on an SP network $PoA_{sc} \leq n$. This is the content of Theorem 2. Second, we show that for every symmetric network topology G that is not an SP network, there exists a game that is played on G for which the PoA can be arbitrarily high. This part is the content of Theorem 2.

Theorem 2. *Let Δ be a feasible CCS game with an underlying graph G. If G is an SP graph then $PoA_{sc}(\Delta) \leq n$.*

In order to complete the characterization it remains to show that for every non-SP network G, there exists a symmetric CCS game on G that has an unbounded price of anarchy.

Theorem 3. *Let G be a non-SP symmetric network. Then, there exists a symmetric CCS game on G for which the price of anarchy is arbitrarily high.*

In order to prove the last theorem, we use the following result, established by Milchtaich [21].

Lemma 1. *[21] A symmetric network G is an SP network if and only if the symmetric network in Fig. 3.1(b) is not embedded in G.*

The network topology in the last lemma is precisely the network topology with the unbounded PoA that motivated our study. The last lemma asserts that this graph topology is embedded in every non-SP network. Thus, in order to establish the assertion of Theorem 1, it remains to show that the unbounded PoA given in Proposition 1 can be *extended* to every network topology that embeds it. This is established in the following lemma.

Lemma 2. *Let G be a symmetric network that is not PoA bounded with respect to sum-cost for a family of symmetric CCS games $\mathcal{F}$, and suppose G is embedded in a symmetric network G'. Then, G' is not PoA bounded with respect to sum-cost for the family $\mathcal{F}$ either.*

For the case of parallel-edge networks, we show that the PoA cannot exceed the maximum edge capacity in the network.

Theorem 4. *Let Δ be a feasible CCS game with an underlying graph G that consists of parallel edge. Let C_m denote the maximum capacity of any edge in G. It holds that $PoA_{sc}(\Delta) \leq C_m$.*

3.2 Price of stability (PoS)

As mentioned above, for uncapacitated symmetric games, $PoS = 1$. In capacitated game, however, the PoS need not be optimal. Moreover, suboptimality is obtained already in parallel-link networks.

Theorem 5. *There exists a symmetric CCS game in which the PoS with respect to sum-cost is $H(n)$.*

Proof. Consider a CCS game with n agents played on a graph that consists of $n+1$ parallel links, $e_1, \ldots, e_{n+1}$, such that for $i \in [n]$, $p_i = 1/i$ and $c_i = 1$; and $p_{n+1} = 1 + \epsilon$ and $c_{n+1} = n$. It is easy to verify that the optimal solution is achieved when all the agents share edge e_{n+1}. However, this profile is not a NE since a single agent can benefit by deviating to edge e_n, incurring a cost of $1/n$ instead of $(1+\epsilon)/n$. Following similar reasonings, agents will continue to deviate, one by one, until reaching the profile in which for every agent $i \in [n]$, agent i uses edge e_i. The cost of this profile is $H(n)$; the assertion follows.

As established in [3], the potential function method can be used to show that the last bound is tight. The proof uses the potential function $\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{x_e(S)}$, and follows the same reasoning as in the uncapacitated case.

Theorem 6. *[3] For every feasible symmetric CCS game, it holds that $PoS_{sc} \leq H(n)$.*

4 The max-cost objective function

In this section we study the max-cost objective function.

4.1 Price of anarchy (PoA)

We first observe that the PoA can be arbitrarily high also with respect to the max-cost function.

Proposition 2. *The PoA with respect to max-cost in CCS games can be arbitrarily high.*

As in the sum-cost case, we wish to characterize network topologies in which the PoA cannot exceed n. Interestingly, we obtain the exact same characterization as in the sum-cost case.

Theorem 7. *A symmetric network topology G is PoA bounded w.r.t. max-cost if and only if G is an SP network.*

For the case of parallel-edge networks, we show that the PoA cannot exceed the maximum edge capacity in the network.

Theorem 8. *Let Δ be a feasible CCS game with an underlying graph G that consists of parallel edge. Let C_m denote the maximum cost of any edge in G. It holds that $PoA_{mc}(\Delta) \leq C_m$.*

4.2 Price of stability (PoS)

For SP graphs, it follows directly from Theorem 7 that the PoS is bounded by n (since PoS is always bounded by PoA). This bound is tight, as follows from the example given in the proof of Theorem 5 . In this example, the unique NE is one in which every agent uses a distinct path, and the maximal cost incurred by any agent is 1, compared to $1/n$ in the optimal solution. For general networks, we establish the following bound.

Theorem 9. *For every CCS game Δ, it holds that $PoS_{mc}(\Delta)$ is bounded by $nH(n)$.*

Proof. Consider the function $\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{x_e(S)}$. It is shown by [3] that this is an exact potential function for the game; i.e., it emulates the change in the cost of a deviating agent. It is easy to verify that for every profile T,

$$sc(T) \leq \Phi(T) \leq H(n) \cdot sc(T). \qquad (2)$$

Let S^* be an optimal solution with respect to max-cost, and consider a NE S that is obtained by running best-response dynamics with an initial profile S^*. We get that $mc(S) \leq sc(S) \leq \Phi(S) \leq \Phi(S^*) \leq H(n)sc(S^*) \leq nH(n)mc(S^*)$, where the second and fourth inequalities follow from Equation 2, the third inequality follows from the fact that Φ is a potential function and S is obtained from S^* through best-response steps, and the last inequality follows from the definition of max-cost. It follows that $mc(S)/mc(S^*) \leq nH(n)$, as promised.

5 Convergence rate of BRD

In this section we study the convergence rate of best-response dynamics (BRD) to a NE. While BRD may in general take exponential number of steps depending on the number of agents to converge [3], the following proposition establishes that in the case of a symmetric, undirected graph, BRD converges to a pure NE within at most n steps, and this is tight. The intuition for this observation is that, in the uncapacitated version, after an agent deviates to some path P (as its best-response), the cost incurred by an agent using this path in the next iteration can only decrease; therefore, P remains a best-response move until all agents converge to the same path.

Observation 2. *For every uncapacitated cost-sharing game, every BRD converges to a NE within at most n steps, independent of the initial profile.*

In contrast, the following proposition shows that the convergence process of a capacitated game may be longer. In particular, we establish a lower bound of $\Omega(n^{3/2})$, even for parallel-link graphs.

Proposition 3. *There exists a symmetric CCS game and a best-response dynamics with convergence time of $\Omega(n^{3/2})$.*

6 Discussion

In this work we introduce a model of capacitated network design games, and study the implications of edge capacities on the existence and quality of Nash equilibria with respect to different objective functions, as well as on the convergence rate of best-response dynamics. We find that the consideration of edge capacities has a significant effect on all the above properties. Our main contribution is a full characterization of network topologies that have a bounded price of anarchy, independent of the edge capacities and costs. Our results suggest many avenues for future research. A few obvious directions include closing the gap of the PoS with respect to the max-cost objective for general networks, the consideration of non-symmetric networks and a better understanding of the convergence rate of best-response dynamics.

References

1. Susanne Albers. On the value of coordination in network design. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 294–303, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
2. Nir Andelman, Michal Feldman, and Yishay Mansour. Strong Price of Anarchy. In *SODA'07*, 2007.
3. Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, Washington, DC, USA, 2004. IEEE Computer Society.
4. Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. In *STOC*, pages 511–520, 2003.
5. Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1230, September 2000.
6. Jacomo Corbo and David Parkes. The price of selfish behavior in bilateral network formation. In *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, PODC '05, pages 99–107, New York, NY, USA, 2005. ACM.
7. Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. In *In Proc. of ACM EC*, pages 108–114, 2003.

8. Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 84–92, New York, NY, USA, 2007. ACM.
9. Amir Epstein, Michal Feldman, and Yishay Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66(1):115–125, May 2009.
10. Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. Convergence time to nash equilibria. In *ICALP*, pages 502–513, 2003.
11. Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, STOC '04, pages 604–612, New York, NY, USA, 2004. ACM.
12. M. Feldman and T. Tamir. Convergence rate of best response dynamics in scheduling games with conflicting congestion effects. *working paper*, 2011.
13. Dimitris Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory Comput. Syst.*, 47(1):113–136, 2010.
14. Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.
15. Ron Holzman and Nissan Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21(1-2):85–101, October 1997.
16. Ron Holzman and Law-yone (Lev-tov) Nissan. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46(2):193–205, 2003.
17. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *in Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
18. Ho lin Chen and Tim Roughgarden. Network design with weighted players. In *In Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architextures (SPAA)*, pages 29–38, 2006.
19. Igal Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Mathematics of Operations Research*, 30:225–244, 2005.
20. Igal Milchtaich. The equilibrium existence problem in finite network congestion games. In *In Proc. of the 2nd Int. Workshop on Internet and Network Economics (WINE)*, pages 87–98, 2006.
21. Igal Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57(2):321–346, 2006.
22. Dov Monderer. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
23. Christos Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, STOC '01, pages 749–753, New York, NY, USA, 2001. ACM.
24. R W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
25. Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, March 2002.